Sites: examples and planning

Marco Mambelli – marco@hep.uchicago.edu

August 9 2011

OSG Summer Workshop

TTU - Lubbock, TX

Open Science Grid

THE UNIVERSITY OF CHICAGO

Summary

- Definition, Communities
- Plan, example
- Some ideas
  - System administrator
  - Who to talk
  - Infrastructure
- Components
- Options for the cluster, topologies
- Batch system, Storage systems
- Grid services

OSG Sites

- Provide one or more of the following capabilities:
  - access to local computational resources using a batch queue
  - interactive access to local computational resources
  - storage of large amounts of data using a distributed file system
  - access to external computing resources on the Grid
  - the ability to transfer large datasets to and from the Grid
- Offer computing resources and data to fellow grid users

General Goals for OSG Sites

- Goal: Build a Robust Infrastructure
  - Consider physical and logical topologies
  - Provide alternate paths when feasible
  - Tune, test, monitor and manage
- Meta-Goal: Protect Services while Maintaining Performance
  - Services should be configured in such a way that they “fail gracefully” rather than crashing. Potentially many ways to do this
  - Tune, test, monitor and manage (as always)

Many communities

- Campus
  - Local resources and support
- Virtual Organization (science experiment)
  - Direction and requirements
- OSG
  - Tools and solutions for Grid computing
  - Specific components (common needs)
  - Ideas from a wider community
  - Support

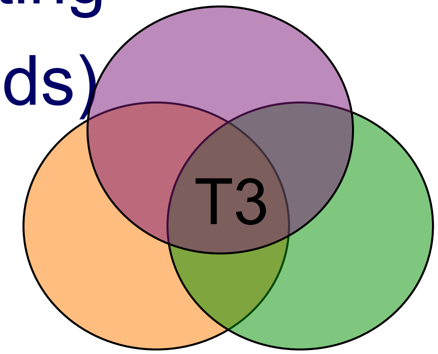

Plan

- Start with requirements from local community and scientific collaboration (VO)
- Define what you need
- See what you can leverage
- Consider the infrastructure
- Plan for expansion
- Have experts checking your plans
- Purchase needed resources
- Install and maintain

Example of Tier 3 definition - ATLAS

- Classify Tier 3s depending on the level of interaction with the Grid and the VO
- Example of activities for a T3g:
  - Run Athena jobs interactively on small data samples.
  - Submit jobs to Grid using pathena (or prun) and retrieve the output
  - Get substantial amount (several TB) ATLAS data to a local storage and keep them.
  - Analyze, using athena or root, a large (TB) data sets in a short time (~1day) in an local batch system
  - Generate and reconstruct Monte Carlo samples locally.
  - Run root jobs interactively for final steps of the analysis.

Example of Tier 3 definition -CMS

- University cluster, not Tier 1-2, unique site design depending on local needs, resources and capabilities
- CMS prefers Tier 3s with full grid access: you can run CRAB jobs on it in exactly the same way as on a Tier 2 or 1
  - Easier to verify a proper setup, e.g. by running identical jobs on the same datasets at Tier 3 vs Tier 2
  - Will provide confidence to physics analysis groups when approving official results
- The USCMS Tier 3 hypernews is an excellent resource for asking questions about what to do in your special situation.

Select a site manager

- A site requires an initial setup effort and a smaller continuous maintenance effort
- If you can get experienced help at your institute, you should do so.
- A person with a clear responsibilities for the cluster is needed. It cannot be a group responsibility.
- Assign one person from your group and a backup, to the site setup effort. If at all possible, the same persons should be responsible for site administration when site is operational.
- Having a backup person will be important. Although the maintenance tasks is envisioned to be light, some of these will have to be done daily or weekly, or it may not be able to wait until the admin returns. Think about rotation of responsibilities after a while.

People you need to know

- Department or university system administrator
- Infrastructure managers: Space, Power and Cooling
- Responsible for networking on campus
- Campus computer security officer
- VO collaborators: e.g. “nearest” Tier 2 and Tier 1
- Hardware representative: group agreements from your campus or your VO
- VO support and coordination (e.g. US CMS or US ATLAS T3 coordination)

OSG role in initial contacts

- Help you to identify your needs
- Provide with vocabulary and examples that make it easier to interact with local experts
- Put you in touch with experts
- Provide intermediaries that can act on your behalf or on your side specially for the initial contact
- Share experiences and solutions

Some notes about infrastructure

- Critical to a well functioning cluster
- Examples of Infrastructure include:
  - Physical space and associated hardware (Racks)
  - Electrical Power and Cooling
  - Networking
  - Computer security / data security
  - System administration and maintenance

Physical space

- Prior to making your computer purchases determine where you will put your hardware
- Keep in mind:
  - 1 Rack of computers is heavy > 1000 lbs
  - Rack of computers is noisy and generates a lot of heat
- Does your University department have a computer room that you can use part of?
- Do you have space for eventual expansion?
- Do you have easy access to machines for repairs?
- Are there costs involved?

Electrical power

- What type of electrical power is available? (110 or 220 V) How much current? (number of circuits)
- Each Dell R710 (used by LHC T3) draws 300W (max) 200W (nominal). i.e. 10 servers in a rack will draw 3000W
- Consider other equipment as well. E.g. UPS.
- Check the load with local safety. Usually 50-70% of the total circuit capacity can be assigned
- Will you have to pay for electricity?

Cooling

- Sufficient cooling important to operation of your cluster
- Some AC systems require a steady load
- Some reference cluster used by LHC T3:
  1. 23U, 927lbs (837 HEPSPEC, 72 cores, 48 TB) - storage on worker nodes – 4745 W (@220V) ~ 16000 BTU/hr ~ 1.4 tons AC (1 ton AC = 12000 BTU/hr; 1 W ~ 3.4 BTU/hr)
  2. 27U, 1279lbs- (837 HEPSPEC, 72 cores, 96 TB) - storage on worker node + extra centralized storage – 5245 W ~17800 BTU/hr ~ 1.5 tons of AC
- Will you have to pay for cooling?

Networking

- OSG sites can benefit from 10GE LANs
  - CPU to storage performance is a potential blocker for large data-size jobs
  - Current hardware makes this feasible
- WAN connectivity also very important
  - OSG is a distributed infrastructure
  - Access to remote resources and data an integral component
  - Generally 1 Gbps or better WAN should be in place…10+ Gbps if you can afford it

Network topologies

- Intranet (private) + Extranet (public)
  - More independent (IP addresses)
  - Manage dual homed hosts

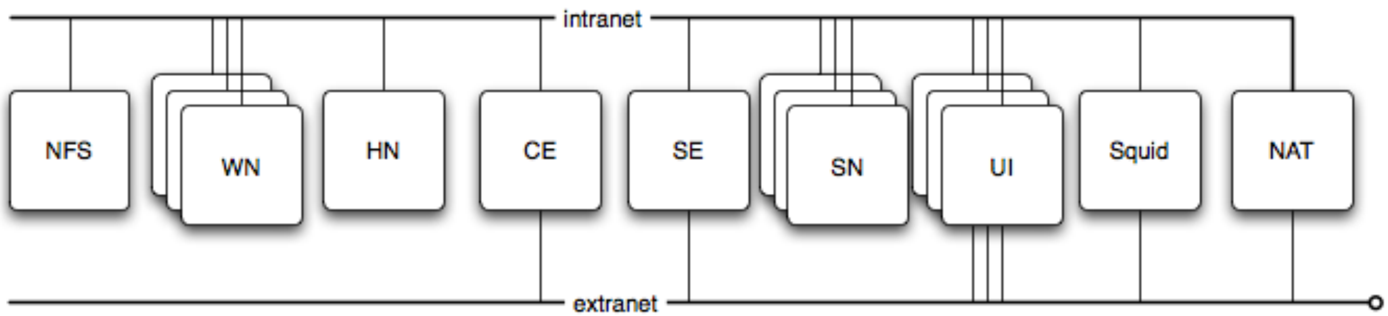

- All on one network

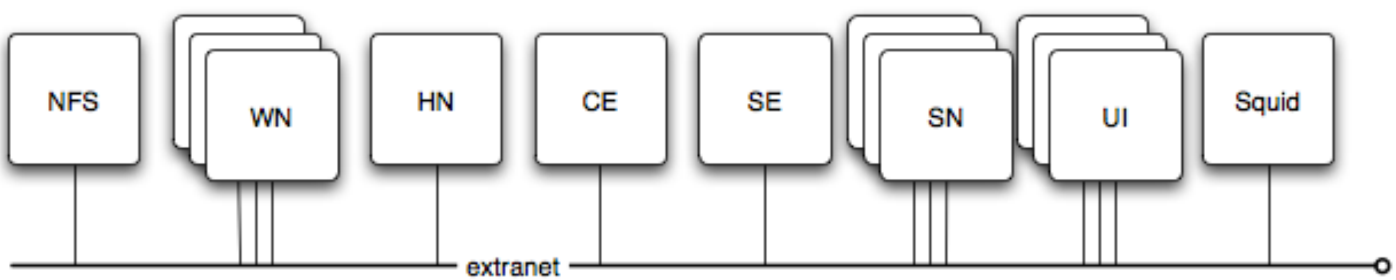

Connection to the public internet

- Connected directly to the internet
- Behind a firewall
  - Cluster more protected
  - Requests to open ports for outgoing or incoming connections
  - May have bandwidth limitations

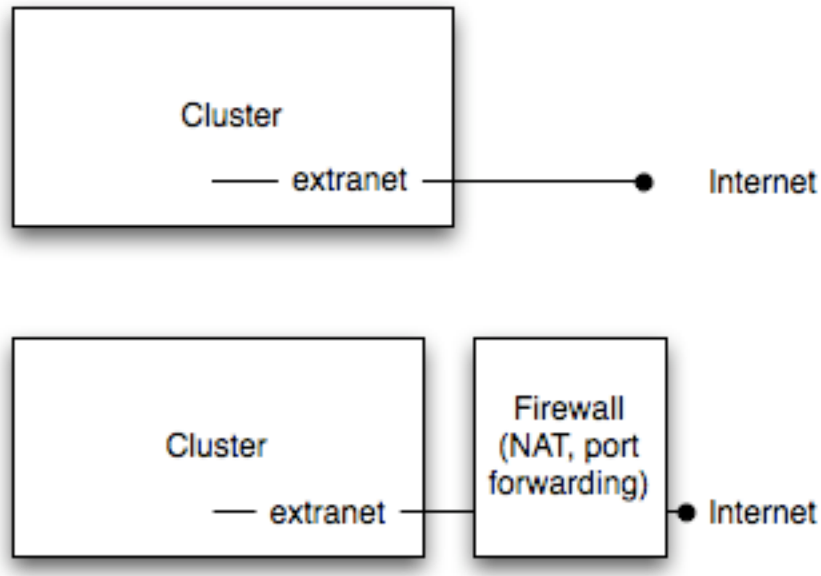

Networking questions

- How far are you from the edge of campus?
- Determine the available bandwidth between your computers and campus backbone?
- Determine the available bandwidth across the campus backbone?
- Determine the available campus bandwidth to Internet 2?
- Is the amount of available bandwidth sufficient for your needs? (100 MB/s ~ 1 TB /day)
- Determine how much networking infrastructure you will have to purchase? Can you use a specific brand of switches? Does your campus require Cisco or another vendor?
- Will you have to pay for bandwidth used?

Networking questions (cont)

- How many public IP address can you get?
- What is the campus firewall policy?
- Some places have several networks
  - Public/restricted to the campus or department
  - Open/limited, e.g. behind a firewall or a traffic shaper
- Do you need a private network for your cluster?
  - Tier 3 examples have public and private networks
  - Added complexity with advantages
  - "No" firewall on private network.

Security

- Who is the department computer security contact? Meet with them.
- Secure computers are vital to our ability to produce science results.
- What are your campus/department computer security policies?
- What will be your role for your cluster?
- We do not want to be the weak link in the computer security chain. - Computer security should not be ignored
- You will hear more in Anand’s talk later

System administration

- How centralized is your campus?
- Who is responsible for machine installation? And for up keep (hardware and software)?
- Does you department have system administrators who can help you?
- Can they administer the machines (OS/ accounts etc)?
- Will you have to do it all but they provide expert guidance?
- What is your data preservation plan? What is your backup strategy

Site Infrastructures

- There are a number of areas to examine where we can add robustness (usually at the cost of **$** or **complexity**!)
  - Networking
    - Physical and logical connectivity
  - Storage
    - Physical and logical connectivity
    - File systems, OS, Software, Services
  - Compute
    - Logical processors, virtualization, local I/O
- Using known configuration or tapping into other people experience may simplify the task (local, VO, OSG)

Storage example

- Increase robustness for storage by providing resiliency at various levels:
  - Network: Bonding (e.g. 802.3ad)
  - Raid/SCSI redundant cabling, multipathing
  - iSCSI (with redundant connections)
  - Disk choices: SATA, SAS, SSD ?
  - **Single-Host resiliency**: redundant power, mirrored memory, RAID OS disks, multipath controllers
  - **Clustered/failover** storage servers
  - Distributed file systems (Hadoop, Xrootd, …): multiple copies, multiple write locations

Site Disk Choices

- Currently a number of disks choices:
  - SATA – Inexpensive, varying RPMs, good BW
  - SAS - Higher quality, faster interface, more $
  - SSDs – Expensive, great IOPs, interface varies
- Reliability is now very good for most choices. NL-SAS (SATA/SAS hybrid) very robust. Range of throughputs and IOPS.
- SSDs have IOPS in the 10K+ region versus fast SAS disks at ~200.   SSD I/O bandwidths usually x2-3 rotating disk

SSDs for Targeted Apps

- SSDs make good sense if IOPS are critical.
  - DB applications are a prime example
  - Server hot-spots (NFS locations, other)
  - Example: Intel SSDs at AGLT2 decreased some operation times from 4 hours to 45 minutes
- New SSDs very robust compared to previous generations…capable of 8+ Petabytes of writes; 5 year lifetimes
- Lower power; bandwidths to 550 MB/sec
- Choice of interface: SATA, SAS, bus-attached. Hitachi/Toshiba have 6 Gbps/SAS

Machines in a site

- Service nodes (as needed):
  - NFS, Cluster Monitoring, Cluster Management
  - Batch Management, User Management, User authentication, CE
  - Data Gateway/buffer, Web data buffer, Data Management, SE
- Interactive nodes (one or more):
  - User login, interactive analysis, submission to local batch and Grid.
  - local user storage area.
- Batch nodes (one or more—two or more for a meaningful batch system):
  - Parallel batch processing queues
  - Storage space for data
- Dedicated storage nodes
- For a very light installation, you can consider an interactive only cluster (nodes may run services and store data)

E.g. US-ATLAS T3g machines

- 2 Service nodes:
  - 1 server for: NFS, Data Gateway/buffer, Cluster Monitoring, Cluster Management
  - 1 server for: Batch Management, Data Management, User Management, Web data buffer
- Interactive nodes (one or more):
  - User login, interactive analysis, submission to local batch and Grid.
  - local user storage area.
- Batch nodes (one or more—two or more for a meaningful batch system):
  - Parallel batch processing queues.
  - Storage space for data.
- Depending on you needs you might add
- Storage nodes for data.
- For a very light installation, you can consider an interactive only cluster.
- Service nodes, in this case will most likely only 1 server (not all services will be needed) or even be a part of an interactive node.

E.g. CMS $100k Tier 3

- Assumptions: (Rob Snihur, Ian Fisk 2009)
  - 6 physicists, (1.4 + 1) TB each
  - Process sample in 24 hrs → 16 nodes w/ 8 cores each
  - Flush & update sample in 12 hrs → 600 Mb/s networking
- Upgradeable RAID chassis ($33k)
- 16 worker nodes ($41k)
- 24-port Gigabit switch ($12k)
- 3 server nodes ($9k)
- Racks and infrastructure ($5k)

OSG documentation

- Reference documents on the TWiki in ReleaseDocumentstion
- Tier3 Web on the TWiki
- Covers mainly Grid services
- Some suggestions on other components of a Tier 3
- User contributions

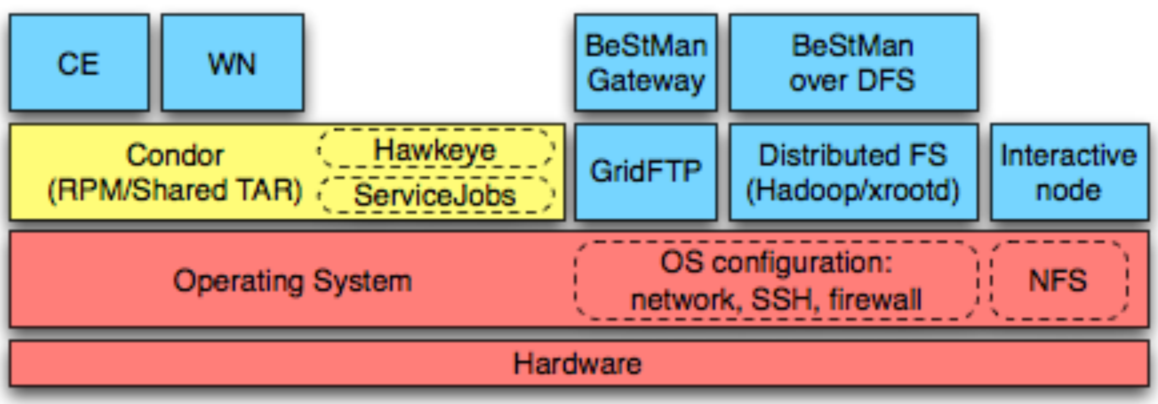

Some notes about the OS

- VDT supported platforms (for the services) http://vdt.cs.wisc.edu/releases/2.0.0/requirements.html
- RHEL 5 (and variants) supported for native pkg.
- RHEL 5 (and variants) most common platform
- Don't forget time synchronization (NTP) https://twiki.grid.iu.edu/bin/view/Tier3/ClusterTimeSetup
- More notes in the Tier 3 documentation (phase 1) https://twiki.grid.iu.edu/bin/view/Tier3/ModulesIntro

Cluster components

- Shared File System (general purpose)
  - Usually NFSv3/NFSv4 (safer) Server. Using NFS to create a shared file system is the easiest way to set up and maintain a Tier 3. NFS is likely to cause performance issues. Read about NFS tuning/setup. Lustre or other commercial solutions (GPFS) can be an alternative.
- Condor Batch Queue
  - A batch queue system is strongly recommended for Tier 3s. Tier 3 documents only provide the installation of Condor (selected because it is one of the most familiar internally to the OSG and hence easily supported by the OSG), but other systems can be used and may be preferable, for example if there is local expertise available in another batch queuing system. The general OSG documentation provides some help for different systems.
- Distributed File System (for data)
  - An optional capability that can be helpful for moving efficiently VO data and other files across the worker nodes. It may also provide data-locality performance improvements to scientific applications. This document covers the installation of Xrootd, a DFS optimized for ROOT files used in the HEP community, although other systems may be used. Tomorrow you will see Hadoop.

Cluster components (cont)

- Cluster deployment and configuration management
  - Provides automatic deployment and some management operations.
  - Some CMS sites use Rocks, a free cluster management solution based on a 'clean reinstall' model. Modifications to the default distribution are done by editing xml files in an admin-friendly way, creating a new distribution, then reinstalling the compute/worker nodes. Rocks has one head node which serves the distribution and numerous other services and which does not require regular reinstall. Nodes can be broken into groups of Rocks appliances, where each group is served a different distribution according to your directions. The 'clean reinstall' model can be convenient for guaranteeing system integrity as compute nodes are working on a very clean and well-defined system. It can pose difficulties when compute nodes are used as interactive nodes or when a particular appliance requires frequent updates, necessitating frequent reinstalls.
  - Some ATLAS sites use Puppet, Ruby-based, is a declarative language for expressing system configuration, a client and server for distributing it, and a library for realizing the configuration. It is more a configuration management tool and allows changes without reinstalling the system.
  - Bcfg2, Cfengine, Cobbler, Modules, Perceus/Warewulf and Quattor are other alternative systems.
- Monitoring
  - Ganglia, Cacti and Nagios are useful monitoring and alarm tools.
- http://www.linuxjournal.com/magazine/taming-beast?page=0,0

Site services in OSG

- Authentication Service
  - enables grid users to authenticate with your site using their grid or voms proxies
- Compute Element
  - enables grid users to run jobs on your site
- Worker Node Client
  - enables grid jobs running on worker nodes to access grid tools
- Storage Element
  - enables grid users to store large amounts of data at your site
- VO Management Service
  - provides functionality for VO Managers to manage the membership information of their users

Authentication Service

- enables grid users to authenticate with your site using their grid or VOMS proxies
- Alternatives:
  - grid-mapfile (edg-mkgridmap)
    - a simple program that contacts VOMS servers and creates a grid-map file
    - easy to install and maintain
    - does not support voms proxies
  - GUMS
    - a web service providing sophisticated controls of how users authenticate
    - supports voms proxies (groups, roles)
    - requires Tomcat to be run as a web service

Compute Element

- enables grid users to run jobs on your site
- Services
  - GRAM: Globus service for job submission
  - GridFTP: grid authenticated file transfer
- Optional services
  - GRAM-WS: web service implementation of GRAM
  - Squid: caching Web proxy
- Managed-fork jobmanager

Worker Node Client

- enables grid jobs running on worker nodes to access grid tools
- Options
  - Shared or local installation
- On worker nodes and compute element
- Required – if supported on your platform

Storage Element

- enables grid users to store or access large amounts of data at your site
- Interfaces
  - Storage Resource Manager (SRM)
    - Dynamic Space Management Support (dyn)
    - No or Static Space Management Support (st)
  - GridFTP
- Doug's talk later today and https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/StorageInfrastructureSoftware
- Several options

Storage Requirements	Min Hardware Requirements	OSG SE Solution
SRM interface, Dynamic Space Management Support	Server with local disk	BeStMan-fullmode
SRM interface, No or Static Space Management Support	Server with local disk or NFS	BeStMan-gateway
SRM interface, No or Static Space Management Support, jobs need root protocol to access data	Multiple servers(>3)	BeStMan-gateway/Xrood
SRM interface, No or Static Space Management Support, file replication	Multiple servers(>4)	BeStMan-gateway/HDFS
SRM interface, Dynamic Space Management Support, file replication, interface to tape backend	Multiple servers (>5)	dcache

Compute and Storage Element

- provide dedicated hardware for the Compute and the Storage Element
- use as many CPUcores and main memory as possible
- avoid running other grid services such as GUMS on the Compute and the Storage Element
- avoid running a file server on the Compute and the Storage Element

VO Management Service

- provides functionality for VO Managers to manage the membership information of their users
- Each VO needs to provide one VOMS
- You need it only if you manage a VO
- If you need groups you can request them to your VO

A very compact OSG site

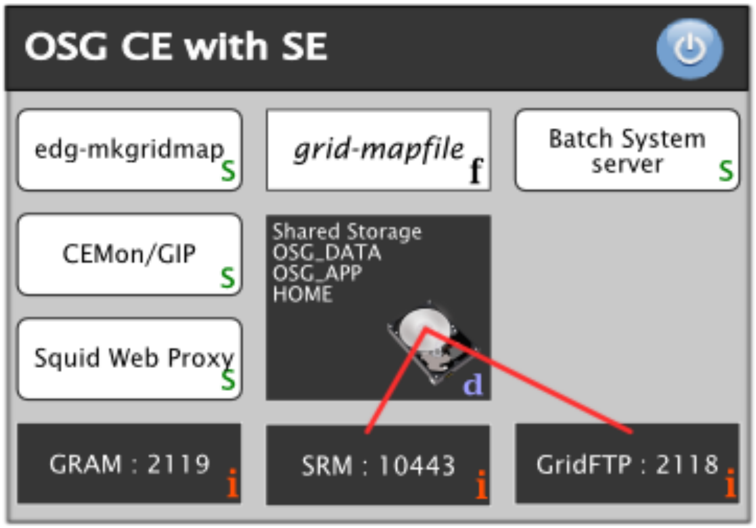

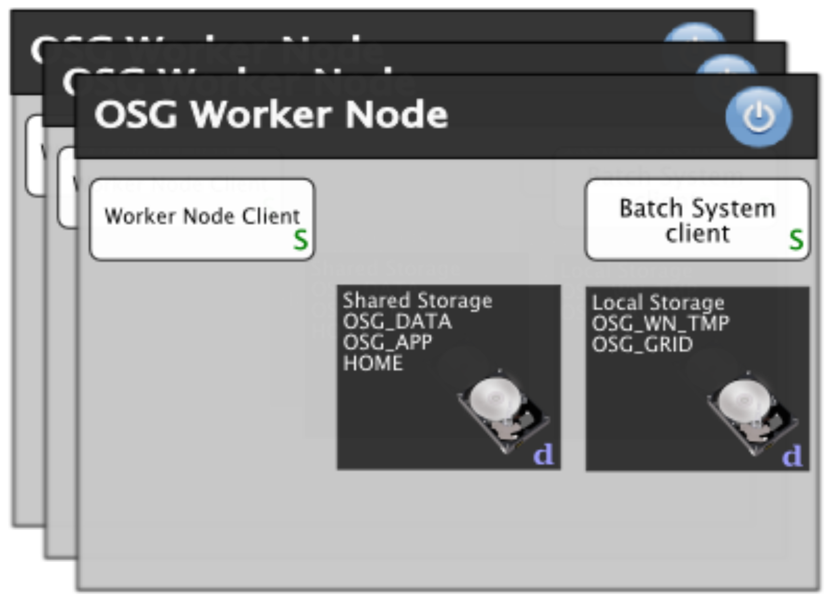

OSG site

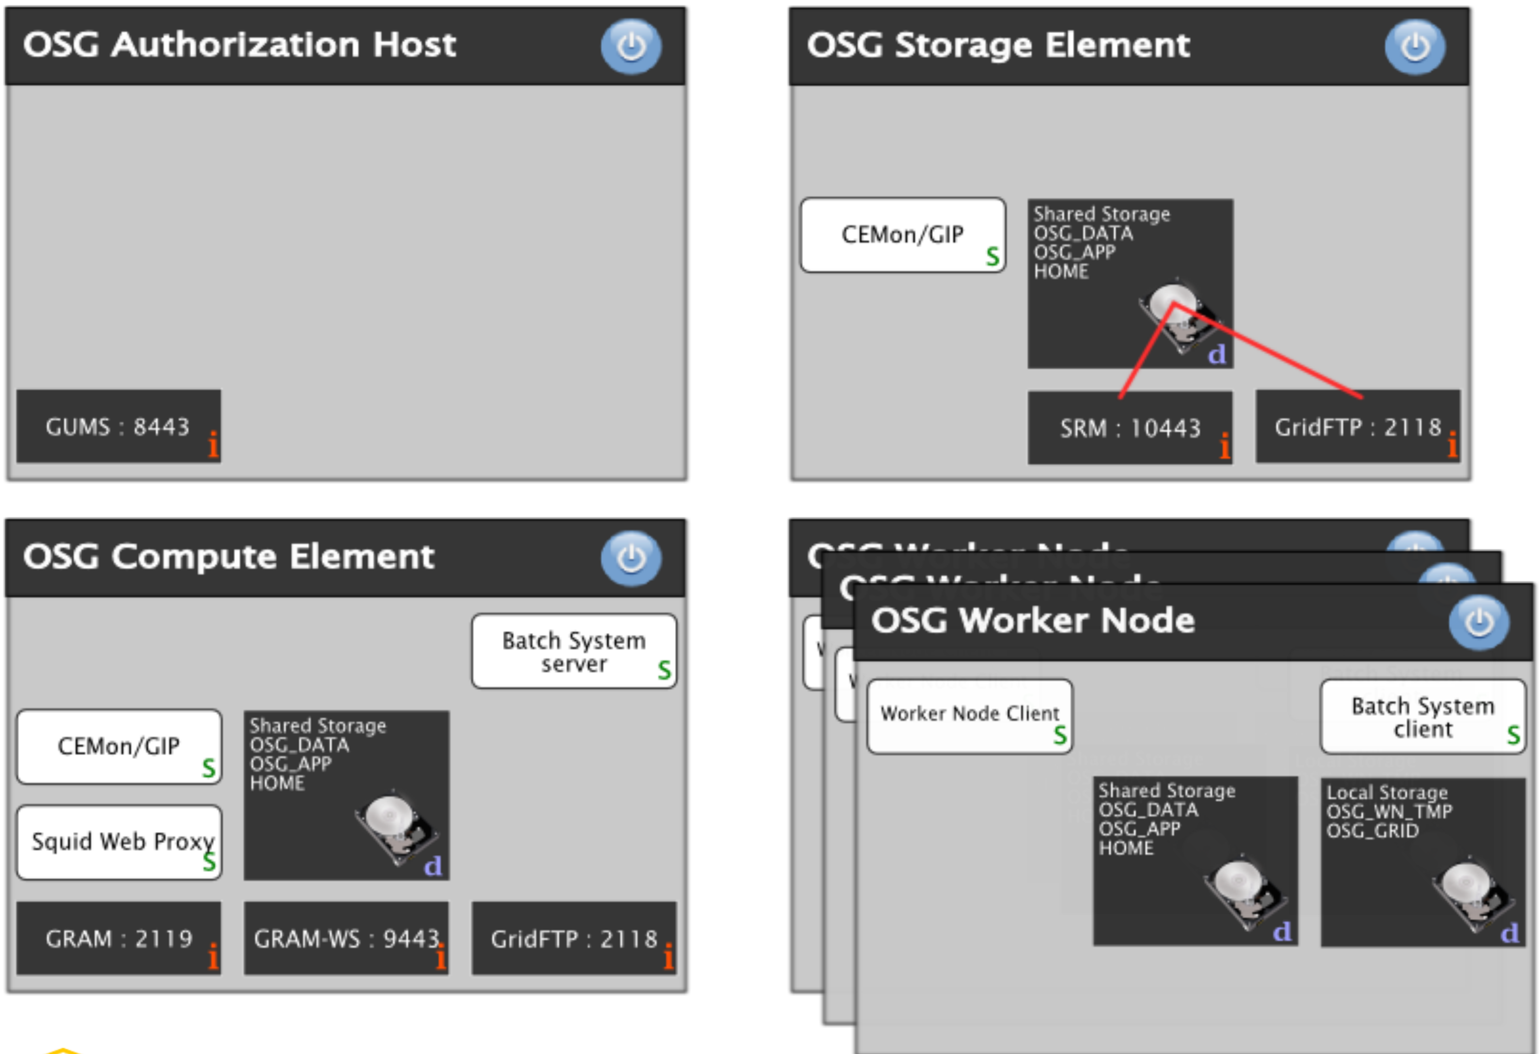
OSG Authorization Host
GUMS : 8443
OSG Storage Element
CEMon/GIP
Shared Storage
OSG_DATA
OSG_APP
HOME
SRM : 10443
GridFTP : 2118
OSG Compute Element
Batch System
server
CEMon/GIP
Squid Web Proxy
Shared Storage
OSG_DATA
OSG_APP
HOME
GRAM : 2119
GRAM-WS : 9443
GridFTP : 2118
OSG Worker Node
Worker Node Client
Batch System
client
Shared Storage
OSG_DATA
OSG_APP
HOME
Local Storage
OSG_WN_TMP
OSG_GRID

Virtual machines

- There are many services in OSG
- VM Allow to install services on separate hosts
  - Isolation and easy load balancing
  - No conflicts for ports or other resources
  - Standard configuration
- This can help also:
  - Backing up critical services
  - Increasing availability, reliability
  - Easing management
- Some penalties
  - Installation and management
  - IO penalty
  - Slower response time (e.g. Xrootd redirector)

General considerations

- Sometimes the additional complexity to add "resiliency" actually decreases availability compared to doing nothing!
- Having test equipment to experiment with is critical for trying new options
- Often you need to trade-off cost vs performance vs reliability (pick 2 ☺)
- Documentation, issue tracking and version control systems are your friends!

Credits

- Thank you to Doug and Rik from the ATLAS Tier 3 coordination
- Thank you to Rob Snihur from the CMS Tier 3 coordination
- Thank you to Shawn McKee (University of Michigan – PI on the ATLAS Great Lakes Tier-2 center)
- Thank you to people writing OSG documentation, specially Site planning Guide

Summary

- There are many components and complex interactions possible in our sites.
- We need to understand our options (frequently site specific) to help create robust, high-performing infrastructures
- Hardware and software choices can give resiliency and performance. Need to tweak for each site.
- Test changes to make sure they actually do what you want (and not something you don't want!) and document what you did
- We are not alone: campus, VO, OSG

Useful links

- Release documentation
  - https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation
- RPM documentation (temporary)
  - https://twiki.grid.iu.edu/bin/view/Documentation/RPMTempDocuments
- Tier 3 documents
  - https://twiki.grid.iu.edu/bin/view/Tier3/WebBook
- OSG Cluster Building Blocks… - OSGAH11
  - https://indico.fnal.gov/contributionDisplay.py?contribId=7&confId=3627
- OSG Site planning guide
  - https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/SitePlanning
- Workshop public TWiki:
  - https://twiki.grid.iu.edu/bin/view/SiteCoordination/SummerWorkshopTutorials2011

?

!